arXiv:1403.2319v2 [cs.LO] 2 Oct 2014

Speeding Up Logico-Numerical Strategy Iteration (extended version)*

David Monniaux[1] and Peter Schrammel[2]

[1] CNRS / VERIMAG
[2] University of Oxford

Abstract. We introduce an efficient combination of polyhedral analysis and predicate partitioning. Template polyhedral analysis abstracts numerical variables inside a program by one polyhedron per control location, with *a priori* fixed directions for the faces. The strongest inductive invariant in such an abstract domain may be computed by upward strategy iteration. If the transition relation includes disjunctions and existential quantifiers (a succinct representation for an exponential set of paths), this invariant can be computed by a combination of strategy iteration and satisfiability modulo theory (SMT) solving. Unfortunately, the above approaches lead to unacceptable space and time costs if applied to a program whose control states have been partitioned according to predicates. We therefore propose a modification of the strategy iteration algorithm where the strategies are stored succinctly, and the linear programs to be solved at each iteration step are simplified according to an equivalence relation. We have implemented the technique in a prototype tool and we demonstrate on a series of examples that the approach performs significantly better than previous strategy iteration techniques.

Keywords: Static analysis, abstract interpretation, strategy iteration, predicate abstraction

1 Introduction

Program verification for unbounded execution times generally relies on finding inductive loop (or procedure) invariants. In the *abstract interpretation* approach, loop invariants are automatically searched within a class known as an *abstract domain*. When dealing with numerical variables, it is common to search for invariants shaped as products of intervals (constraints $l \leq x \leq u$ with the program variable x and bounds l, u), convex polyhedra (constraint system $\boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{c}$ with the matrix $\boldsymbol{A}$, the vector of program variables $\boldsymbol{x}$, and the vector of bounds $\boldsymbol{c}$),

* The research leading to these results has received funding from the European Research Council under the European Unions Seventh Framework Programme (FP7/20072013) / ERC grant agreement 306595 "STATOR" and the ARTEMIS Joint Undertaking under grant agreement number 295311 "VeTeSS"

or restricted classes of convex polyhedra such as *octagons* (constraints $\pm x_i \pm x_j \leq c$ with program variables x_i, x_j and a bound c). Intervals and octagons are instances of *template polyhedra*: polyhedra where $\boldsymbol{A}$ is fixed *a priori*, whereas in the general polyhedral approach, $\boldsymbol{A}$ is discovered. The restriction to fixed $\boldsymbol{A}$ reduces the problem to finding suitable values for a fixed number of bounds $\boldsymbol{c}$, and even, for certain classes of transitions, minimizing these bounds using *strategy iteration* [1] (also known as *policy iteration*) or other techniques, thereby producing the least (or *strongest*) inductive invariant in the abstract domain. In contrast, for unknown $\boldsymbol{A}$ the number of constraints (rows in the $\boldsymbol{A}$ matrix) may grow quickly and is most often limited by *widening* heuristics [2].

One common weakness of all these abstract domains is that they represent only *convex* numerical properties; it is for instance impossible to represent $|x| \geq 1$ where $|x|$ denotes the absolute value of x. An analyzer running on

```
while(...) {
  ...
  if (abs(x) >= 10) { assert(x != 0); ... }
}
```

will flag a warning on the assertion, because at this point the non-convex property $x \leq -10 \vee x \geq 10$ has been abstracted away. In order to relieve this weakness, some recent approaches [1,3,4] advocate convex abstractions only at a *cut-set* of all program locations — a subset such that removing all points in this subset cuts all cycles in the control-flow graph, e.g. all loop heads within a structured program. In between such distinguished locations, program executions are exactly represented (or at least represented more faithfully) as solutions to satisfiability modulo theory (SMT) formulas. This is equivalent to replacing the original control flow graph by a multigraph whose vertices are the distinguished nodes and the edges are the simple paths between the distinguished nodes.

Variants on this basic idea include combinations with invariant inference with widenings [3, 4] and with strategy iteration [1]. With such approaches, the test **if** (abs(x) >= 10) is interpreted as the disjunction $x \leq -10 \vee x \geq 10$ and the code is analyzed in both contexts $x \leq -10$ and $x \geq 10$. If n tests are used in succession, the number of cases to analyze may be 2^n, but such methods eschew exhaustive enumeration for as-needed consideration of paths inside the code through SMT-solving.

There still remains a difficulty: what if disjunctive invariants are needed at the distinguished nodes? Assume for example, that predicate abstraction is used to handle programs that contain constructs other than linear arithmetic, for example, pointers, dynamic data stuctures, non-linear arithmetic, etc. A similar situation arises in *reactive programs* for control applications, where a main loop updates global variables at each iteration, including Booleans (or, more generally, variables belonging to a finite enumerated type) encoding "modes" of operation. Such a system has a single distinguished control point (the head of the main loop), yet, one wants to distinguish invariants according to the mode of operation of the system. Assuming modes are defined by the values of the n Boolean

variables, this can be achieved by splitting the loop head into 2^n distinct control nodes and computing one invariant for each of them.

Should we apply a max-strategy iteration modulo SMT algorithm [1] to these 2^n control nodes, its running time would be in the worst case proportional to 2^{d2^n} where d is the number of disjuncts in the formula defining the semantics of the program. Worse, it would construct linear programming problems with $2^n\ell$ unknowns, where ℓ is the number of rows in the $\boldsymbol{A}$ matrix. While high worst case complexity is not necessarily an objection (many algorithms behave in practice better than their worst case), constructing exponentially-*sized* linear programs at every iteration of the algorithm is certainly too costly. We thus previously left this partitioning variant as an open problem [1, §9] [5, §3.5].

Contributions. The main contribution of this paper is an algorithm that computes the *same* result as these prohibitively expensive methods proposed in [1,6], but limits the costs by computing on-the-fly a form of equivalence between constraint bounds (of which there are exponentially many) and constructing problems whose size depends on the number of these equivalence classes. These equivalence classes, in intuitive terms, distinguish Boolean variables *with respect to the abstraction chosen* (the $\boldsymbol{A}$ matrix). This is a novel aspect that distinguishes our approach from quotienting techniques (*e.g.* [7]). In contrast to [5] that uses approximations to scale, we aim at computing the strongest invariant. Finally, we show the results of an experimental evaluation conducted with our prototype implementation that demonstrate the largely improved performance in comparison to previous strategy iteration techniques.

2 Strategy Iteration Basics

Let us now recall the framework of strategy iteration over template linear constraint domains [6], reformulating it to the setting of programs with linear arithmetic and Boolean variables. As explained above, Boolean variables may be introduced by predicate abstraction or by the encoding of the control flow as in reactive systems. Similar to [1], this allows us to represent an exponential number of paths as a single compact transition formula. We then explain why previous algorithms [1, 6] have unacceptable complexity when instantiated on our exponential number of control nodes.

Notation. We shall often talk both about formal variables appearing in logical formulas and about individual values they may take, particularly those obtained as satisfying instances of formulas; if needed, we shall distinguish values by denoting them $\hat{x}, \hat{y} \ldots$ as opposed to variables $x, y, \ldots$. Variables $x_1, \ldots, x_n$ are denoted collectively as a vector $\boldsymbol{x}$. When discussing satisfaction of logical formulas, we shall note $(\boldsymbol{x}, \boldsymbol{y}) \models F$ to mean explicitly that $\boldsymbol{x}, \boldsymbol{y}$ are free variables of F that should satisfy F.

2.1 Program Model and Abstract Domain

We model a program as a transition system with m rational variables $\boldsymbol{x} = (x_1, \ldots, x_m) \in \mathbb{Q}^m$ (the *numeric state*) and n Boolean variables $\boldsymbol{b} = (b_1, \ldots, b_n) \in \mathbb{B}^n$ (the *Boolean state*), where $\mathbb{B} = \{0, 1\}$. Let $I = (b_1^0, \ldots, b_n^0, x_1^0, \ldots, x_m^0) = (\boldsymbol{b}^0, \boldsymbol{x}^0)$ be the initial state. The transition relation τ is of the form $\exists y_1, \ldots, y_E \in \mathbb{Q},\ \exists p_1, \ldots, p_d \in \mathbb{B}.\ T$ where T is a quantifier-free formula in negation normal form, whose atoms are either propositional (b_i, $\neg b_i$, p_i, $\neg p_i$), linear (in)equalities ($\sum \alpha_i x_i + \sum \alpha'_i x'_i + \sum \beta_i y_i \bowtie c$, where the α_i, α'_i, β_i and c are rational constants) with $\bowtie \in \{\leq, <, =\}$, and y_i variables to encode nondeterminism;[3] the free variables of τ are $x_1, \ldots, x_m, b_1, \ldots, b_n, x'_1, \ldots, x'_m, b'_1, \ldots, b'_n$ where primed (respectively, unprimed) variables denote the state after (respectively, before) the transition. Furthermore, we add p_i variables to each disjunction with non-propositional literals, *i.e.*, $x \leq 3 \vee x \geq 6$ becomes $(p_i \wedge x \leq 3) \vee (\neg p_i \wedge x \geq 6)$. This encoding is necessary to uniquely identify each disjunct by a Boolean proposition and extract it from a SAT model. The free variables of T are thus grouped into $\boldsymbol{b}, \boldsymbol{b}', \boldsymbol{x}, \boldsymbol{x}', \boldsymbol{p}, \boldsymbol{y}$ where $(\boldsymbol{b}, \boldsymbol{x})$ and $(\boldsymbol{b}', \boldsymbol{x}')$ define respectively the departure and arrival states and $\boldsymbol{p}, \boldsymbol{y}$ stand for intermediate values and choices.

Example. We consider the following running example (a variant of the classical thermostat model):

```
1  bool error = 0, heat_on = 1;
2  bool fan_on = read_button();
3  real t = 16;
4  while(1) {
5    real te = read_external_temp();
6    assume(14<=te && te<=19);
7    fan_on = read_button() ? !fan_on : fan_on;
8    if(!error && (t<15 || t>30)) error = 1;
9    else if(heat_on && t>22) heat_on = 0;
10   else if(!error && heat_on && t<=22) t = (15*t + te)/16 + 1;
11   else if(!error && !heat_on && t<18) heat_on = 1;
12   else if(!heat_on && t>=18) t = (15*t + te)/16;
13 }
```

This program has the following transition relation T with $n = 3$ Boolean variables $\boldsymbol{b} = (e, h, f)$ (short for (error,heat_on,fan_on)), $m = 1$ numerical variables $\boldsymbol{x} = (t)$, $d = 3$ path choice variables $\boldsymbol{p} = (p_0, p_1, p_2)$), $\boldsymbol{y} = (te)$, and initial states $\neg e \wedge h \wedge t = 16$:

$$
\begin{array}{ll}
\neg p_0 \wedge p_1 \wedge p_2 \wedge \neg e \wedge e' \wedge (h = h') \wedge & t > 30 \wedge t' = t \vee \\
\neg p_0 \wedge p_1 \wedge \neg p_2 \wedge \neg e \wedge e' \wedge (h = h') \wedge & t < 15 \wedge t' = t \vee \\
p_0 \wedge p_1 \wedge p_2 \wedge h \wedge \neg h' \wedge (e = e') \wedge & 22 < t \leq 30 \wedge t' = t \vee \\
p_0 \wedge p_1 \wedge \neg p_2 \wedge \neg e \wedge h \wedge \neg e' \wedge h' \wedge & t \leq 22 \wedge 14 \leq te \leq 19 \wedge t' = \frac{15t+te}{16} + 1 \vee \\
p_0 \wedge \neg p_1 \wedge p_2 \wedge \neg e \wedge \neg h \wedge \neg e' \wedge h' \wedge & 15 \leq t < 18 \wedge t' = t \vee \\
p_0 \wedge \neg p_1 \wedge \neg p_2 \wedge \neg h \wedge \neg h' \wedge (e = e') \wedge & t \geq 18 \wedge 14 \leq te \leq 19 \wedge t' = \frac{15t+te}{16}
\end{array}
$$

[3] This limitation to linear (in)equalities may be lifted using, *e.g.*, *linearization* techniques [8]. For integer values, simple transformations should be performed, *e.g.*, $x < y$ to $x \leq y - 1$. Floating-point operations may also be relaxed to nondeterministic real operations [9, §4.5].

The disjuncts stem from lines 9 –13; line 9 produces two path choices.

Abstract Domain. Let $\boldsymbol{A}$ be a $\ell \times m$ rational matrix, with rows $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_\ell$.[4] An element ρ of the abstract domain $D^\sharp$ is a function $\mathbb{B}^n \to \overline{\mathbb{Q}}^\ell$ with $\overline{\mathbb{Q}} = \mathbb{Q} \cup \{-\infty, +\infty\}$. $\rho(\boldsymbol{b}) = \boldsymbol{c}$ means that at a Boolean state $\boldsymbol{b}$ the vector of numerical variables $\boldsymbol{x}$ is such that $\boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{c}$ coordinate-wise. Moreover, we write $\rho(\boldsymbol{b}) = -\boldsymbol{\infty}$ if any coordinate $c_i = -\infty$, meaning that the Boolean state $\boldsymbol{b}$ is unreachable (because $\boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{c}$ is *false*). We note $\gamma(\rho)$ the set of states $(\boldsymbol{b}, \boldsymbol{x})$ verifying these conditions. $\mathbb{Q}^\ell$ is ordered by coordinate-wise $\leq$, inducing a point-wise ordering $\sqsubseteq$ on $D^\sharp$. γ is thus monotone w.r.t. $\sqsubseteq$ and the inclusion ordering on sets of states; note that it is not injective in general. We denote by $\rho(i, \boldsymbol{b})$ the i-th coordinate of $\rho(\boldsymbol{b})$. ρ is said to be an *inductive invariant* if it contains the initial state ($\boldsymbol{A}\boldsymbol{x}^0 \leq \rho(\boldsymbol{b}^0)$) and it is stable by transitions:

$$\forall \boldsymbol{b}, \boldsymbol{x}, \boldsymbol{b}', \boldsymbol{x}'.\ (\boldsymbol{b}, \boldsymbol{x}) \in \gamma(\rho) \wedge (\boldsymbol{b}, \boldsymbol{x}, \boldsymbol{b}', \boldsymbol{x}') \models \tau \implies (\boldsymbol{b}', \boldsymbol{x}') \in \gamma(\rho'),$$

otherwise said $\forall \boldsymbol{b}, \boldsymbol{x}, \boldsymbol{b}', \boldsymbol{x}'.\ \rho(\boldsymbol{b}) \neq \bot \wedge \boldsymbol{A}\boldsymbol{x} \leq \rho(\boldsymbol{b}) \wedge (\boldsymbol{b}, \boldsymbol{x}, \boldsymbol{b}', \boldsymbol{x}') \models \tau$

$$\implies \rho(\boldsymbol{b}') \neq \bot \wedge \boldsymbol{A}\boldsymbol{x}' \leq \rho(\boldsymbol{b}').$$

The main contribution of this paper is an effective way to compute the least inductive invariant ρ in this abstract domain with respect to inclusion ordering.

2.2 Strategy Iteration

Recall that the original strategy iteration algorithm [6] applies to disjunctive systems of linear inequalities (of exponential size in d) induced by the collecting semantics of the program over template polyhedra (see Equ. (1) below). Previous work [1] improves the algorithm by keeping the system implicit, only extracting a linear size system at any given time using SMT solving. Note that τ, after replacing each free Boolean variable by a Boolean constant, is equivalent to a disjunction of (at most) 2^d formulas of the form $\exists \boldsymbol{y}.C$ where C is a conjunction of non-strict linear inequalities, and d is the number of Boolean existential quantifiers in τ. In both algorithms, a *strategy*[5] selects one disjunct in C for each template row i'. Hence, we can use these algorithms in our setting by selecting a disjunct for each template row *and* each Boolean valuation for $(\boldsymbol{b}', \boldsymbol{p})$. This motivates the following definition:

A *strategy* associates with each Boolean state $\boldsymbol{b}'$ and each constraint index $1 \leq i' \leq \ell$ either the special value $\bot$, meaning that $\boldsymbol{b}'$ is unreachable and denoted by $\sigma(i', \boldsymbol{b}') = \pi(i', \boldsymbol{b}') = \bot$, or a pair $\big(\sigma(i', \boldsymbol{b}'), \pi(i', \boldsymbol{b}')\big)$ where $\sigma(i', \boldsymbol{b}') \in \mathbb{B}^n$ is a Boolean state and $\pi(i', \boldsymbol{b}') \in \mathbb{B}^d$ gives "path choices". Once $\pi \in \mathbb{B}^d$ is chosen, the result of substituting $T[\pi/\boldsymbol{p}]$ is a conjunction of linear inequalities and a

[4] One can also make $\boldsymbol{A}$ depend on $b_1, \ldots, b_n$ so as to apply a non-uniform abstraction, adding minor complication to algorithms and proofs. We chose to describe uniform abstraction for the sake of brevity and clarity.

[5] The word "strategy" (or "policy") arises from an analogy between the system of min-max + monotone affine linear equalities whose least solution yields the least invariant in the domain [1, 6] and the system of min-max + barycentric inequalities whose least solution is the value of a two-player Markov game.

propositional formula (in the variables $\boldsymbol{b}, \boldsymbol{b}'$); let T_π denote the conjunction of these linear inequalities.[6]

Algorithm. Let us now see how the algorithm iterates until the least inductive invariant is reached. The algorithm maintains, at iteration number k, a current strategy (σ_k, π_k). Initially, the abstract value ρ_0 is $\bot$ everywhere save at the initial Boolean state $\boldsymbol{b}^0$, where $\rho(\boldsymbol{b}^0) = \boldsymbol{A}\boldsymbol{x}^0$; σ_0 and π_0 are $\bot$ everywhere. For $k \geq 0$, the strategy yields ρ_{k+1} as the least fixed point $\mu_{\geq \rho_k} \Psi_\pi$ greater than ρ_k, Ψ_π being an order-continuous operator on the lattice $(\mathbb{B}^n \times \{1, \ldots, \ell\}) \to \overline{\mathbb{Q}}$ defined as:

$$\Psi_\pi(\rho) \stackrel{\triangle}{=} (i', \boldsymbol{b}') \mapsto \sup \left\{ \boldsymbol{A}_i \boldsymbol{x}' \mid \exists \boldsymbol{x}, \boldsymbol{y}.\ T_{\pi(i', \boldsymbol{b}')} \wedge (\boldsymbol{A}\boldsymbol{x} \leq \rho(\sigma(i', \boldsymbol{b}'))) \right\} \quad (1)$$

We explain in §2.3 how to compute this fixed point; let us now see how σ_{k+1} and π_{k+1} are obtained from σ_k and π_k, and how termination is decided [1, §6.5]. Each iteration goes as follows: for all Boolean states $\hat{\boldsymbol{b}'} \in \mathbb{B}^n$ and all $\hat{\boldsymbol{b}}$ with $\rho_k(\hat{\boldsymbol{b}}) \neq -\boldsymbol{\infty}$:

1. construct formula $T[\hat{\boldsymbol{b}}/\boldsymbol{b}, \hat{\boldsymbol{b}'}/\boldsymbol{b}']$, that is, T where variables $\boldsymbol{b}$ and $\boldsymbol{b}'$ have been replaced by Boolean values $\hat{\boldsymbol{b}}$ and $\hat{\boldsymbol{b}'}$;
2. conjoin it with constraints $\boldsymbol{A}\boldsymbol{x} \leq \rho_k(\boldsymbol{b})$ and $\boldsymbol{A}_i\boldsymbol{x} > \rho_k(i, \boldsymbol{b}')$, thus obtaining
$$T[\hat{\boldsymbol{b}}/\boldsymbol{b}, \hat{\boldsymbol{b}'}/\boldsymbol{b}'] \wedge \boldsymbol{A}\boldsymbol{x} \leq \rho_k(\hat{\boldsymbol{b}}) \wedge \boldsymbol{A}_i\boldsymbol{x}' > \rho_k(i, \hat{\boldsymbol{b}'}) \quad (2)$$
3. check whether this formula (in free variables $x_1, \ldots, x_m, x'_1, \ldots, x'_m, y_1, \ldots, y_E$, $p_1, \ldots, p_d$) is satisfiable;
4. if this formula is satisfiable, ρ_k does not describe an inductive invariant: the satisfying instance describes a transition from a state $(\boldsymbol{b}, \boldsymbol{x})$ to a state $(\boldsymbol{b}', \boldsymbol{x}')$ such that $(\boldsymbol{b}, \boldsymbol{x})$ lies within the invariant but $(\boldsymbol{b}', \boldsymbol{x}')$ does not; this solution yields a new strategy $\pi_{k+1}(i, \hat{\boldsymbol{b}'}) = \hat{\boldsymbol{p}}$ and $\sigma_{k+1}(i, \hat{\boldsymbol{b}'}) = \hat{\boldsymbol{b}}$, which *improves* on the preceding one [1, §6.3];
5. if $\pi_{k+1}(i, \hat{\boldsymbol{b}'})$ and $\sigma_{k+1}(i, \hat{\boldsymbol{b}'})$ are not set by the preceding step, leave them to the their previous values $\pi_k(i, \hat{\boldsymbol{b}'})$ and $\sigma_k(i, \hat{\boldsymbol{b}'})$; if none have been updated, this means ρ_k is the least inductive invariant, thus **exit**;
6. otherwise, compute $\rho_{k+1} = \mu_{\geq \rho_k} \Psi_{\pi_{k+1}}$ and continue iterating.

The main loop of this algorithm enumerates each of the $2^{(n+d)\ell 2^n}$ strategies at most once. Remark the important improvement condition: at every iteration but the last, $\Psi_\pi(\rho) > \rho$. Since there is a finite number of strategies that may deem ρ non-inductive and each of them is chosen at most once, we are guaranteed to terminate with the least fixed point (without using any widening) within a finite number of steps.

Example. Let us analyze our running example using the box template $(t, -t)^T$. Assume the current abstract value[7] $\rho_k(i, \bar{e}hf) = 16$ for $i \in \{1, 2\}$. To compute

[6] This conjunction corresponds to a *merge-simple statement* of [1] or to a *path* of [3,4].
[7] For better readability, we write, for example, $\bar{e}hf$ for the value $(0, 1, 1)$ of (e, h, f).

an improved strategy σ_{k+1}, π_{k+1}, we have to check all values of $(\hat{\boldsymbol{b}}, \hat{\boldsymbol{b}'})$, *e.g.*, $(\bar{e}hf, \bar{e}'h'f')$: instantiating Equ. 2 with these values (with T from our running example and $i = 1$) gives

$$\big(p_0 \wedge p_1 \wedge \neg p_2 \wedge t \leq 22 \wedge 14 \leq te \leq 19 \wedge (t' = \tfrac{15t+te}{16} + 1)\big) \wedge \big(t = 16\big) \wedge \big(t' > 16\big)$$

which is satisfied, for instance, by the model $(\hat{\boldsymbol{x}}, \hat{\boldsymbol{x}}', \hat{\boldsymbol{p}}) = (16, 17, (1, 1, 0))$. Hence, we update the strategy by setting $\sigma_{k+1}(1, \bar{e}hf) = \bar{e}hf$ and $\pi_{k+1}(1, \bar{e}hf) = (1, 1, 0)$, which induces

$$T_{\pi_{k+1}(1,\bar{e}hf)} = (t \leq 22 \wedge 14 \leq te \leq 19 \wedge (t' = \tfrac{15t+te}{16} + 1))\,.$$

After having checked all $(\hat{\boldsymbol{b}}, \hat{\boldsymbol{b}'})$, we can compute the strategy value, *i.e.*, the fixed point of $\Psi_{\pi_{k+1}}$, which updates $\rho_{k+1}(1, \bar{e}hf)$ to $\frac{365}{16}$ in this case.[8] The way this is computed is explained in the next section.

2.3 Computing the Strategy Value

We recall now how to compute the strategy fixed point $\mu_{\geq\rho}\Psi_\pi$ [1, §6.4], under the condition that $\Psi_\pi(\rho) \geq \rho$ (which is always the case, because of the way π is chosen).

The first step is to identify the Boolean states $\boldsymbol{b}$ "abstractly unreachable": such $\boldsymbol{b}$ form the least set Z containing all $\boldsymbol{b} \neq \boldsymbol{b}^0$ such that $\pi(i, \boldsymbol{b}) = \bot$ and stable by: if $\boldsymbol{b}' \neq \boldsymbol{b}^0$ is such that $\sigma(i, \boldsymbol{b}') \in Z$ then $\boldsymbol{b}' \in Z$; for all $\boldsymbol{b} \in Z$, set $\rho(\boldsymbol{b}) := -\boldsymbol{\infty}$.

Construct a system of linear inequalities in the unknowns $v_{i,\boldsymbol{b}}$ for $\boldsymbol{b} \in \mathbb{B}^n$ and $1 \leq i \leq \ell$, plus fresh variables: for all $\boldsymbol{b}' \notin Z$, for all $1 \leq i' \leq \ell$ such that $\rho(i', \boldsymbol{b}') < +\infty$, add the inequalities

- $\boldsymbol{A}_j \boldsymbol{x} \leq v_{j,\sigma(i',\boldsymbol{b}')}$ for all $1 \leq j \leq \ell$ ("in departure state invariant")
- $\boldsymbol{A}_{i'} \boldsymbol{x}' \geq v_{i',\boldsymbol{b}'}$ ("bounded by arrival state")
- those from the conjunction $T_{\pi(i',\boldsymbol{b}')}$ ("follows the transition relation")

where variables $\boldsymbol{x}$ and $\boldsymbol{y}$ have been replaced by fresh variables (each different $i, \boldsymbol{b}'$ has its own set of fresh replacements). $\rho(i, \boldsymbol{b}')$ is obtained by linear programming as the maximum of $v_{i',\boldsymbol{b}'}$ satisfying this system. This linear program has solutions, otherwise the strategy σ, π would not have been chosen; if it has no *optimal* solution it means that $\rho(i', \boldsymbol{b}') = +\infty$.

Note that these $O(2^n \ell)$ linear programs have $O((2\ell + E)2^n)$ variables and a system of inequalities of size $O(2^n |T|)$ where $|T|$ is the size of formula T. It is in fact possible to replace these $O(2^n \ell)$ linear programs by two linear programs of size $O((2\ell + E)2^n)$: first, one using the ∞-abstraction (see [10, §8,9]) to obtain which of the $v_{i',\boldsymbol{b}'}$ go to $+\infty$, then another for maximizing $\sum v_{i',\boldsymbol{b}'}$ restricted to the $v_{i',\boldsymbol{b}'}$ found not to be $+\infty$ by the ∞-abstraction.

3 Our Algorithm

Notice three difficulties in the preceding algorithm: there are, *a priori*, 2^{2n} SMT-solving tests to be performed at each iteration; the linear programs have expo-

[8] By maximizing t for any te, we get $\frac{15 \cdot 22 + 19}{16} + 1 = \frac{365}{16}$.

nential size; and there are at most $2^{(n+d)\ell 2^n}$ strategies, thus a doubly exponential bound on the number of iterations. In intuitive terms, the first two difficulties stem from the explicit expansion of the exponential set of Boolean states, despite the implicit representation of the exponential set of execution paths between any two control (Boolean) states $\boldsymbol{b}$ and $\boldsymbol{b}'$, a weakness that we shall now remedy.

3.1 Strategy Improvement Step

The first difficulty is the easiest to solve: the 2^{2n} SMT-tests, one for each pair $(\boldsymbol{b}, \boldsymbol{b}')$ of control states, can be folded into one single test where the $\boldsymbol{b}$ and $\boldsymbol{b}'$ also are unknowns to be solved for.

Note that the structure of ρ, $\mathbb{B}^n \to \overline{\mathbb{Q}}^\ell$, can be viewed as $\{1, \ldots, \ell\} \to (\mathbb{B}^n \to \overline{\mathbb{Q}})$. Hence, we need not store a $2^n \times \ell$ array of rationals (or infinities), but we can implement it efficiently as an array (of size ℓ) of MTBDDs [11] with the bounds $c_{i,j}$ in the leaves. Assume for a given template row i, we have s_i different bounds $c_{i,j}$, and denote $\phi_{i,j}$ the propositional formula describing the set of Boolean states that map to bound $c_{i,j}$. Then, observe that $\phi_{i,j}$ for $1 \leq j \leq s_i$ form a partition of $\mathbb{B}^n$ (that is, $\bigvee_{j=1}^{s_i} \phi_{i,j}$ is a tautology and each $\phi_{i,j} \wedge \phi_{i,k}$, $j \neq k$, is unsatisfiable). We use the notation $\rho(i) = \{\phi_{i,1} \to c_{i,1}, \ldots, \phi_{i,s_i} \to c_{i,s_i}\}$ to represent an MTBDD, and $\rho(i, \boldsymbol{b}) = c_{i,j}$ to obtain the bound $c_{i,j}$ for state $\boldsymbol{b}$ for template row i.

Strategy improvement condition. In Equ. 2, one may replace $\boldsymbol{A}\boldsymbol{x} \leq \rho(\boldsymbol{b})$ and $\boldsymbol{A}_i\boldsymbol{x} > \rho(i', \boldsymbol{b}')$ respectively by ψ_1 and ψ_2:

$$\psi_1 \stackrel{\triangle}{=} \bigwedge_i \bigvee_{j=1}^{s_i} \phi_{i,j}(\boldsymbol{b}) \wedge \boldsymbol{A}_i\boldsymbol{x} \leq c_{i,j} \tag{3}$$

$$\psi_2 \stackrel{\triangle}{=} \bigvee_{j=1}^{s_i} \phi_{i',j}(\boldsymbol{b}') \wedge \boldsymbol{A}_i\boldsymbol{x}' = c_{i',j} + \Delta \wedge \Delta > 0 \tag{4}$$

Remark that $\psi =_{def} \psi_1 \wedge T \wedge \psi_2$ is satisfiable iff there is a transition from $(\boldsymbol{b}, \boldsymbol{x})$ inside the invariant defined by ρ to $(\boldsymbol{b}', \boldsymbol{x}')$ outside of it. The same applies if we replace T in ψ by a *slicing* or *cone of influence* T_i of T with respect to the value of $\boldsymbol{A}_i\boldsymbol{x}$, that is, a formula T_i such that $(\exists p_1, \ldots, p_d \in \mathbb{B},\ \exists y_1, \ldots, y_E \in \mathbb{Q}.\ T) \wedge \boldsymbol{A}'_i\boldsymbol{x}' \geq v$ and $(\exists p_1, \ldots, p_d \in \mathbb{B},\ \exists y_1, \ldots, y_E \in \mathbb{Q}.\ T_i) \wedge \boldsymbol{A}'_i\boldsymbol{x}' \geq v$ are equivalent (w.r.t $\boldsymbol{b}, \boldsymbol{x}, \boldsymbol{b}', v$). When T is compiled from a program, such a T_i may be obtained using program slicing. The strategy iteration algorithm progresses regardless of Δ as long as $\Delta > 0$. It is however likely that maximizing Δ leads to faster convergence [1, p. 26], because there is no backtracking in max-strategy iteration and hence a *locally optimal*, *i.e.*, greedy, strategy cannot be disadvantageous. Maximizing Δ may be performed by *optimization modulo theory* techniques [12–14].

Obtaining a solution $\boldsymbol{b}, \boldsymbol{x}, \boldsymbol{b}', \boldsymbol{x}', \boldsymbol{y}, \boldsymbol{p} \models \psi$ enables us to improve the strategy by setting $\sigma(i, \boldsymbol{b}') := \boldsymbol{b}$ and $\pi(i, \boldsymbol{b}') := \boldsymbol{p}$, as in §2.2.

Algorithm 1 IMPROVE: Selecting the strategy improvement

1: $stable := true$
2: **for** $i' \in \{1, \ldots, \ell\}$ **do**
3: $U := false$ // U defines the set of $\boldsymbol{b}'$ such that $\pi(i', \boldsymbol{b}')$ has been updated.
4: **while** $\neg U \wedge \left(\bigwedge_i \bigvee_{j=1}^{s_i} \phi_{i,j}(\boldsymbol{b}) \wedge \boldsymbol{A}_i\boldsymbol{x} \leq c_{i,j}\right) \wedge T_{i'} \wedge \left(\bigvee_{j=1}^{s_i} \phi_{i,j}(\boldsymbol{b}') \wedge \boldsymbol{A}_i\boldsymbol{x}' = c_{i,j} + \Delta\right) \wedge \Delta > 0$ is satisfiable **do**
5: $\langle \hat{\boldsymbol{b}}, \hat{\boldsymbol{x}}, \hat{\boldsymbol{b}'}, \hat{\boldsymbol{x}'}, \hat{\boldsymbol{p}}, \hat{\boldsymbol{y}} \rangle :=$ a model of the above formula (optionally of max. Δ)
6: $F := T_i[\hat{\boldsymbol{x}}/\boldsymbol{x}, \hat{\boldsymbol{y}}/\boldsymbol{y}] \wedge \neg U$
7: $stable := false$
8: **while** F is satisfiable **do**
9: $\langle \hat{\boldsymbol{b}_1}, \hat{\boldsymbol{b}'_1}, \hat{\boldsymbol{p}_1} \rangle :=$ a model of F
10: $G := F[\hat{\boldsymbol{b}_1}/\boldsymbol{b}, \hat{\boldsymbol{p}_1}/\boldsymbol{p}]$
11: $F := F \wedge \neg G$
12: $\pi[i', G] := \hat{\boldsymbol{p}}$ // $\pi[i', G] := \hat{\boldsymbol{p}}$ means "in the mapping $\boldsymbol{b}' \mapsto \pi(i', \boldsymbol{b}')$,
13: $\sigma[i', G] := \hat{\boldsymbol{b}}$ // replace all images of $\boldsymbol{b}'$ satisfying formula G
14: $U := U \vee G$ // by $\hat{\boldsymbol{p}}$" (respectively for σ).
15: **end while**
16: **end while**
17: **end for**

Example. Let us assume the following current abstract value in the analysis of our running example:

$$\rho(1) = \{\neg e \wedge h \to 16, \quad e \vee \neg h \to -\infty\}$$
$$\rho(2) = \{\neg e \wedge h \to -16, \; e \vee \neg h \to -\infty\}$$

We build Equ. 2 using Equs. 3 and 4:

$$\psi = T \wedge \begin{pmatrix} (\neg e \wedge h \wedge t \leq 16 & \vee (e \vee \neg h) \wedge t \leq -\infty) \wedge \\ (\neg e \wedge h \wedge -t \leq -16 & \vee (e \vee \neg h) \wedge -t \leq -\infty) \wedge \\ (\neg e' \wedge h' \wedge (t' = 16 + \Delta) & \vee (e' \vee \neg h') \wedge (-t' = -\infty + \Delta)) \end{pmatrix} \wedge \Delta > 0$$

This formula is satisfied, *e.g.*, by the model $(\hat{\boldsymbol{b}}, \hat{\boldsymbol{b}'}, \hat{\boldsymbol{x}}, \hat{\boldsymbol{x}'}, \hat{\boldsymbol{p}}) = (\bar{e}hf, \bar{e}hf, 16, 17, (1, 1, 0))$. Hence, we update $\sigma(1, \bar{e}hf) := \bar{e}hf$ and $\pi(1, \bar{e}hf) := (1, 1, 0)$. We have to repeat this check excluding the above solution to find other models, *e.g.*, $(\bar{e}h\bar{f}, \bar{e}hf, 16, 17, (1, 1, 0))$.

Improving the strategy this way would however be costly, since we would have to do it one $\hat{\boldsymbol{b}'}$ at a time (by naive model enumeration).

Model generalization. There is however a better way by *generalizing* from an obtained model to a set of $\hat{\boldsymbol{b}'}$ that can be updated at once: Notice now that, fixing $\hat{\boldsymbol{x}}$ and $\hat{\boldsymbol{y}}$ arising from a solution, $\psi[\hat{\boldsymbol{x}}/\boldsymbol{x}, \hat{\boldsymbol{y}}/\boldsymbol{y}]$ becomes a purely propositional formula, whose models also yield suitable solutions for $\boldsymbol{b}, \boldsymbol{b}', \boldsymbol{p}$. Fix $\hat{\boldsymbol{b}}$ and $\hat{\boldsymbol{p}}$ from a solution, then the free variables are now only the $\boldsymbol{b}'$; then for any solution $\hat{\boldsymbol{b}'}$ of $\psi[\hat{\boldsymbol{x}}/\boldsymbol{x}, \hat{\boldsymbol{y}}/\boldsymbol{y}, \hat{\boldsymbol{b}}/\boldsymbol{b}, \hat{\boldsymbol{p}}/\boldsymbol{p}]$, we can set $\pi(\hat{\boldsymbol{b}'}, i) := \hat{\boldsymbol{p}}$ and $\sigma(\hat{\boldsymbol{b}'}, i) := \hat{\boldsymbol{b}}$. We can thus improve strategies for whole sets of $\hat{\boldsymbol{b}'}$ at once in nondeterministic systems.

Our strategy improvement algorithm (procedure IMPROVE, Alg. 1) thus proceeds as follows: it maintains a set U of "already improved" values of $\boldsymbol{b}'$, and

Algorithm 2 ITERATE: Main strategy iteration algorithm

for $i \in \{1, \ldots, \ell\}$ **do**
  $\phi_{1,i} := (\boldsymbol{b} = \boldsymbol{b}^0)$; $c_{1,i} := \boldsymbol{A}_i \boldsymbol{x}^0$; $\phi_{2,i} := (\boldsymbol{b} \neq \boldsymbol{b}^0)$; $c_{2,i} := -\infty$
end for
stable := *false*
while $\neg$*stable* **do**
  IMPROVE
  if $\neg$*stable* **then**
    COMPUTE-STRATEGY-VALUE (see §3.2)
  end if
end while

requests $(\boldsymbol{b}, \boldsymbol{b}', \boldsymbol{p})$ by SMT-solving as described above, with the additional constraint that $\boldsymbol{b}' \notin U$; if no such solution is found, it terminates, having done all improvements, otherwise it generalizes $\boldsymbol{b}'$ to a whole set of solutions as described above, and improves the strategy for all these $\boldsymbol{b}'$. The strategy π, σ and the set U are stored in BDDs.

Example. Let us assume we have the current abstract value

$$\rho(1) = \{\neg e \wedge h \to \tfrac{365}{16},\ \ e \vee \neg h \to -\infty\}$$
$$\rho(2) = \{\neg e \wedge h \to -16,\ e \vee \neg h \to -\infty\}.$$

Moreover, assume that we have obtained the model of ψ: $(\hat{\boldsymbol{b}}, \hat{\boldsymbol{b}}', \hat{\boldsymbol{x}}, \hat{\boldsymbol{x}}', \hat{\boldsymbol{p}}) = (\bar{e}hf, \bar{e}\bar{h}f,\ \tfrac{365}{16}, \tfrac{365}{16}, (1,1,1))$. Substituting the values of this solution for $\boldsymbol{x}$ and $\boldsymbol{x}'$ in formula ψ, we get $F = p_0 \wedge p_1 \wedge p_2 \wedge \neg e \wedge h \wedge \neg e' \wedge \neg h'$. Now, we substitute the above values for $\boldsymbol{b}$ and $\boldsymbol{p}$ in F, which gives us $G = \neg e' \wedge \neg h'$. We update the strategy σ, π for the whole set of states satisfying G, *i.e.*, $\{\bar{e}\bar{h}f, \bar{e}\bar{h}\bar{f}\}$ at once, and we add G to U. Then we ask the SAT solver again for a model of the formula $F \wedge \neg G$, which is unsatisfiable in this example. We continue enumerating the solutions of ψ, but this time excluding U, *i.e.*, we call the SMT solver with $\psi \wedge \neg U$, which is unsatisfiable in our example. Hence, we have completed strategy improvement for the first template row. For row 2, we proceed similarly and obtain the same strategy update. The associated strategy value computation yields the abstract value ρ:

$$\rho(1) = \{\neg e \to \tfrac{365}{16},\qquad\qquad e \to -\infty\}$$
$$\rho(2) = \{\neg e \wedge h \to -16,\ \neg e \wedge \neg h \to -22,\ e \to -\infty\}.$$

Lemma 1. IMPROVE *terminates in at most exponential time. At the end, stable is false if and only if the strategy needed updating (otherwise said, $\gamma(\rho_k)$ was not an inductive invariant), in which case σ, π contain the next strategy σ_{k+1}, π_{k+1}.*

Proof. Each iteration of the outer (resp. inner) loop removes at least one solution for $\boldsymbol{b}'$ from $\neg U$ (resp. F), and there are 2^n of them. The updates to π and σ have been explained in the preceding paragraphs.

Theorem 1. ITERATE *(Alg. 2) terminates in at most $2^{(n+d)m2^n}$ iterations, with the final ρ being equal to that computed by the algorithm of §2.2, yielding the least inductive invariant in the domain.*

Proof. The correctness of the result ensues from the correctness of the path-focused strategy iteration approach [1], the correctness of the improvement strategy (Lemma 1) and the correctness of the strategy value computation (proved in §3.2), with the remark that our new algorithm can be considered an instance of the path-focused iteration scheme: the difference with the instance described in §2.2 is that we store ρ in an efficient way and the way we pick the improved strategy, neither of which matters for correctness.

Strategy iteration terminates in at most as many iterations as there are possible strategies: here, for each of the 2^n states $\boldsymbol{b}'$ and i-th constraint ($1 \leq i \leq m$), there are 2^n possible choices for $\sigma(i, \boldsymbol{b}')$ and 2^d choices for $\pi(i, \boldsymbol{b}')$, thus the bound.

3.2 Computing the Strategy Value with Fewer Unknowns

There remains the second difficulty: computing the value of a given strategy, that is, computing $\rho(\boldsymbol{b})$ for $\boldsymbol{b} \in \mathbb{B}^n$, thus solving linear programs with at least $m2^n$ variables [1, §6.4]. We solve this difficulty by remarking that $\rho(i, \boldsymbol{b})$ is the same for all $\boldsymbol{b}$ in the same equivalence class with respect to $\sim_i$: $\boldsymbol{b}_1 \sim_i \boldsymbol{b}_2 \iff \pi(i, \boldsymbol{b}_1) = \pi(i, \boldsymbol{b}_2) \wedge \sigma(i, \boldsymbol{b}_1) = \sigma(i, \boldsymbol{b}_2)$. Assuming $\boldsymbol{b} \mapsto \sigma(i, \boldsymbol{b})$ and $\boldsymbol{b} \mapsto \pi(i, \boldsymbol{b})$ are stored as MtBdds (or, equivalently, n ordinary Bdds for $\boldsymbol{b} \mapsto \sigma(i, \boldsymbol{b})$ and d for $\boldsymbol{b} \mapsto \pi(i, \boldsymbol{b})$, each containing a bit of the image), the equivalence classes are obtained as Bdds using the reverse images of these functions.

We then apply the algorithm from §2.3, but instead of the whole set of $\rho(i, \boldsymbol{b})$ unknowns for $\boldsymbol{b} \in \mathbb{B}^n$ and $1 \leq i \leq m$, we only pick one unknown $c_{i,j}$ per equivalence class; these unknowns define ρ in the form expected by the strategy improvement step of §3.1. Remark that, if the equivalence classes are computed as Bdds, it is trivial to turn them into logical formulas $\phi_{i,j}$ of linear size w.r.t. that of the Bdd. Notice that also the ∞-abstraction technique [10, §8,9] also applies. Let $\bar{\boldsymbol{b}}^i$ denote the equivalence class of $\boldsymbol{b}$ with respect to $\sim_i$; π directly maps from equivalence classes as $\pi(i, \bar{\boldsymbol{b}}^i) \stackrel{\triangle}{=} \pi(i, \boldsymbol{b})$ (resp. for σ).

Example. Let us assume the current abstract value

$$\rho(1) = \{\neg e \to \tfrac{365}{16}, \qquad e \to -\infty\}$$
$$\rho(2) = \{\neg e \wedge h \to -16, \neg e \wedge \neg h \to -22, e \to -\infty\}.$$

Moreover, assume that we have computed the following strategy for the first template row: $\sigma(1, \bar{e}\bar{h}f) = \sigma(1, \bar{e}\bar{h}\bar{f}) \in \{\bar{e}\bar{h}f, \bar{e}\bar{h}\bar{f}\}$ and $\pi(1, \bar{e}\bar{h}f) = \pi(1, \bar{e}\bar{h}\bar{f}) = (1, 0, 0)$. Then the states $\bar{e}\bar{h}f$ and $\bar{e}\bar{h}\bar{f}$ will be in the same equivalence class, because both bounds will have the same value in the strategy fixed point. Hence, we have to generate only one set of constraints for both states when solving the LP problem that characterizes the strategy fixed point ρ.

We finally obtain[9] $\begin{cases} \rho(1) = \{\neg e \to \frac{365}{16}, & e \to -\infty\} \\ \rho(2) = \{\neg e \wedge h \to -16, \neg e \wedge \neg h \to -\frac{71}{4}, & e \to -\infty\}. \end{cases}$ This is actually the strongest inductive abstract invariant of our program: $\neg e \wedge h \wedge 16 \leq t \leq \frac{365}{16} \vee \neg e \wedge \neg h \wedge \frac{71}{4} \leq t \leq \frac{365}{16}$.

[9] By maximizing $-t$ for any te in $t \geq 18 \wedge 14 \leq te \leq 19 \wedge t' = \frac{15t+te}{16}$, we get $-\frac{15 \cdot 18 + 14}{16} = -\frac{71}{4}$.

Theorem 2. *Let $\rho^\sharp$ be the result of the modified strategy evaluation and $\rho_{k+1} = \mu_{\geq \rho_k} \Psi_{\pi_{k+1}}$ be the result of the original strategy evaluation. Then for all $i, \boldsymbol{b}$, $\rho(i, \boldsymbol{b}) = \rho^\sharp(i, \bar{\boldsymbol{b}}^i)$.*

Proof. The original strategy evaluation computes $\rho_{k+1} = \mu_{\geq \rho_k} \Psi_{\pi_{k+1}}$. Remark that ρ_{k+1} is thus the limit of the ascending sequence $u_0 = \rho_k$, $u_{j+1} = \Psi_{\pi_{k+1}}(u_j)$; furthermore, from the definition of $\Psi_{\pi_{k+1}}$ and the form of the equivalence classes, for any $j \geq 1$, $u_j(i, \boldsymbol{b})$ does not depend on the choice of $\boldsymbol{b}$ in an equivalence class of $\sim_i$. It follows that the same limit is obtained by keeping for each j only one $u_j(i, \boldsymbol{b})$ per equivalence class. This corresponds to iterating

$$\Psi^\sharp_\pi(\rho) \stackrel{\triangle}{=} (i', \bar{\boldsymbol{b}'}^{i'}) \mapsto \sup \left\{ \boldsymbol{A}_i \boldsymbol{x}' \mid \exists \boldsymbol{x}, \boldsymbol{y}.\ T_{\pi(i', \boldsymbol{b}')} \wedge (\boldsymbol{A}\boldsymbol{x} \leq \rho(i', \sigma(\bar{\boldsymbol{b}'}^{i'}))) \right\} \tag{5}$$

As in §2.3, the modified strategy evaluation computes the least fixed point $\rho^\sharp$ of $\Psi^\sharp_{\pi_{k+1}}$ greater than $(i, \bar{\boldsymbol{b}}^i) \mapsto u_1(i, \boldsymbol{b})$. But, from the remark above, this implies that for all $(i, \boldsymbol{b})$, $\rho_{k+1}(i, \boldsymbol{b}) = \rho^\sharp(i, \bar{\boldsymbol{b}}^i)$.

3.3 Abstraction Through Limitation of Partitioning

Even though we have taken precautions against unnecessarily large numbers of unknowns by grouping "equivalent" Boolean states together, it is still possible that the number of equivalence classes to consider grows too much as the algorithm proceeds. It is however possible to freeze them permanently, for instance to their last sufficiently small value. Only small modifications to the algorithms are necessary: The strategy value computation (§3.2) remains the same except that the equivalence classes are never recomputed. Let $\phi_{i,1}, \ldots, \phi_{i,s_i}$ denote the propositional formulas (in $\boldsymbol{b}$) defining the equivalence classes with respect to constraint number i. In the strategy improvement step (§3.1) $\sigma(i, j) \in \mathbb{B}^n$ (resp. $\pi(i, j)$) is now defined for the index $1 \leq j \leq s_i$ of an equivalence class with respect to constraint i.

The correctness proofs stay the same, except that instead of computing the least fixed point in $(\{1, \ldots, \ell\} \times \mathbb{B}^n) \to \overline{\mathbb{Q}}$ we compute it in $\left(\bigsqcup_{1,\ldots,\ell} E_i\right) \to \overline{\mathbb{Q}}$ where E_i is the set of equivalence classes associated with constraint i; the latter lattice is included in the former.

3.4 Combination with Predicate Abstraction

We have described so far a method for computing template polyhedral invariants on each element of a partition of the state space according to the value of Booleans $b_1, \ldots, b_n$. These Booleans may be replaced by arbitrary predicates $\chi_1, \ldots, \chi_n$: it suffices to replace T by $T \wedge \bigwedge_i (b_i \Leftrightarrow \chi_i)$.

3.5 Strategy Iteration with Partitioning is EXPTIME-hard

In preceding work without partitioning [1,15], the single-exponential upper bound was shown to be reached by a contrived example program, and the decision

problem associated with the least invariant computation ("given a template, a transition relation, an initial state and a bad state, is there an inductive invariant that excludes the bad state") was shown to be Σ_2^p-complete. We have an NEXPTIME upper bound on the problem. We will now prove EXPTIME-hardness for the problem with partitioning.

Let Π be an EXPTIME problem. Consider a Turing machine $\mathcal{M}$ deciding Π, with a single tape over the alphabet $\{0, 1\}$ with time bounded by $2^{P(n)}$ and finite state in Σ. Let x be an input of size n to $\mathcal{M}$; we are going to describe a program $\mathcal{P}$ of length proportional to $P(n)$ such that its execution would yield the same result as running $\mathcal{M}$ over x. $\mathcal{P}$ only uses Boolean operations for discrete state, and affine linear operations for continuous state.

Let $\mathcal{M}_x$ be the Turing machine $\mathcal{M}$ where x has been substituted into the input; $|\mathcal{M}_x| \simeq |\mathcal{M}| + |x|$. The tape is modeled as a couple of natural integers (l, r) with $2^{P(n)}$ bits, where

- l represents the bits strictly to the left of the read-write head, the $2^{P(n)} - 1$-th bit representing the bit on the tape just left of the read-write head, and the least order bit representing the bit on the tape $2^{P(n)}$ positions left of the head;
- r represents the bits to the right of the read-write head, the $2^{P(n)} - 1$-th bit representing the bit on the tape under the read-write head, and the least order bit representing the bit on the tape $2^{P(n)} - 1$ positions right of the head.

l and r are initialized to 0 (empty tape).

A step of the Turing machine $\mathcal{M}_x$ is simulated as follows:

- the bit b under the read-write head is obtained by taking the $2^{P(n)} - 1$-th bit of r, by comparing r to $K = 2^{2^{P(n)}-1}$;
- the bit w just to the left of the read-write head is similarly obtained by comparing l to K;
- the bit b' to be written to the tape under the head, the direction of movement of the head and the next state are computed according to the rules of $\mathcal{M}$;
- if the tape is to be moved to the left, then a parallel update is made: $l := 2(l - wK)$ and $r := r/2 + b'K$
- if the tape is to be moved to the right, then a parallel update is made: $r := 2(l - bK)$ and $l := l/2 + b'K$

Since there are $2^{P(n)}$ steps to be simulated, we loop over the simulated step using a binary counter S with $P(n)$ bits. This loop ends when the Turing machine under simulation enters a final state.

Note that, in the program, l and r are always natural numbers. This is because, when we execute $l/2$ (resp. $r/2$), the low-order bit of l (resp. r) is necessarily 0, otherwise it would mean that $\mathcal{M}$ would be using more than $2^{P(n)}$ bits of tape. Also, the operations wK, bK, $b'K$ are defined not by non-linear multiplications, but by case analysis over the bits w, b and b'. All resulting elementary operations are thus linear over the reals.

The last issue to solve is how to create K. We cannot write it as a constant in the program, because it has $2^{P(n)}$ bits — the program would have exponential size. Instead, we prepend to the program $K := 1$ followed by a sequence of $2^{P(n)} - 1$ doublings ($K := 2K$), implemented by a loop over binary counter S of length $P(n)$.

We thus obtain a program of length $O(P(n) + n)$ (because of the operations over binary counter S of length $P(n)$). It has $P(n) + 2\lceil \log_2 |\Sigma| \rceil + 3$ bits of discrete storage (not counting control flow: 3 for b, b', w, $P(n)$ for the binary counter S, $2 \log_2 |\Sigma|\rceil$ for implementing the state transition).

Now note the execution of $\mathcal{P}$ is fully deterministic. In addition, along an execution trace (p, S), where p is the control point in $\mathcal{P}$ and S the binary counter, takes distinct values (S is incremented once per loop iteration and p follows control inside the loop). Thus, an interval analysis that has a different interval for l and r for each value of (p, S) will essentially simulate the concrete execution of $\mathcal{P}$ and obtain an *exact* result. Such an interval analysis can thus decide whether $\mathcal{P}$ terminates in "accepting" or "rejecting" answers, and thus whether $\mathcal{M}$ accepts or rejects x.

Despite repeated attempts, we have not yet been able to narrow the interval at proving NEXPTIME-completeness. It is thus possible that worst-case complexity is actually better.

4 Experiments

We have prototypically implemented the algorithm in the static analyzer ReaVer [16] (written in OCaml and taking Lustre code as input) using the LP solver QSOpt_Ex[10], the SMT solver Yices[11] and the BDD package Cudd[12]. The implementation makes heavy use of incremental SMT solving.

Tested variants of the algorithm. We implemented the following variants of the algorithm to compare their performance:

(n) Naive model enumeration using SMT solving per template row as explained in the first part of §3.1. This corresponds to updating π and σ in Alg. 1 using the model obtained in line 5 ($G = (\boldsymbol{b'} = \hat{\boldsymbol{b}}')$) without doing lines 6 to 11 and 15.

(t) Enhancement of (n) by trying to reapply successfully improving models to other template rows.

(s) Symbolic encoding of template rows and model enumeration over the whole template at once, *i.e.*, lines line 2 and 17 are omitted because the template row i' becomes part of the SMT formula to be solved for in line 4, and is then retrieved from the model returned in line 5.

(g) Alg. 1 with *generalization* as described in §3.1, but without the inner iterations (*i.e.*, without lines 7 to 9 and 15) that search for models of the purely

[10] version 2.5.6, `http://www.dii.uchile.cl/~daespino/ESolver_doc/main.html`

[11] version 1.0.40, `http://yices.csl.sri.com/`

[12] version 2.4.2, `http://vlsi.colorado.edu/~fabio/CUDD/`

propositional formula F. Hence, (g) obtains the models to be generalized from the SMT formula in line 4 only.

(m) Alg. 1 as given.

All these variants reduce the number of unknowns in the LP problem using equivalence classes (see §3.2). Furthermore, we used an implementation of the original max-strategy algorithm [6] (GS07), and the improvements using SMT solving proposed in [15] (GM11). Note that these latter two algorithms need to enumerate $\mathcal{O}(2^n)$ control states (where n is the number of Booleans in the recurrent state). Yet, the size of the control flow graphs (see Table 1) generated using the method described in [16, §7.3] is often far smaller than the worst case. The difference between GS07 and GM11 is essentially that, for each template row, the former tests all strategies to find an improvement, whereas the latter asks the SMT solver to find an improving strategy in the disjunction of available strategies.

It is important to note that all these variants of the algorithm return the *same* invariants, *i.e.* the strongest invariants in the domain $\mathbb{B}^n \to A$ where A is a given template abstract domain. The only difference is the way the strategy improvement is computed.

Comparisons. We performed two kinds of comparisons:[13]

1. We evaluated the scalability of various variants of the max-strategy improvement algorithm on small benchmarks (1-, 2-, and 3-dimensional array traversals, parametric in size by duplicating functionality and adding Boolean variables) that exhibit the strategy and state space explosion expected to occur in larger benchmarks. We used box and octagonal templates, giving a total of 96 benchmarks.
2. We compared the max-strategy improvement algorithm with standard forward analysis with widening and with abstract acceleration [17,18] (both using widening after two iterations and applying two descending iterations[14]) on reactive system models (traffic lights [19], our thermostat, car window controller [20], and drug pump [21]), again deriving the more complex variants 2 and 3 by adding and duplicating functionality (e.g. branching *multiple* drug pumps to a patient and checking the concentration in the blood).

Results. The first comparison (see Fig. 1) shows that the various variants of the algorithm behave quite differently in terms of runtime: The GM11 improvement is on average 22% faster than the original algorithm. (t) and (s) scale better than (n). It is interesting to observe that (t) and (s) perform similarly although their algorithms are very different. The most important optimization of the strategy improvement algorithm proposed in this paper is the generalization step which makes it scale several orders of magnitude better than the other variants, because it avoids naive model enumeration. The results indicate that the full Alg. 1

[13] The examples and detailed experimental results can be found on `http://www.cs.ox.ac.uk/people/peter.schrammel/reaver/maxstrat/`.

[14] We did not not observe any improvement in precision beyond these values.

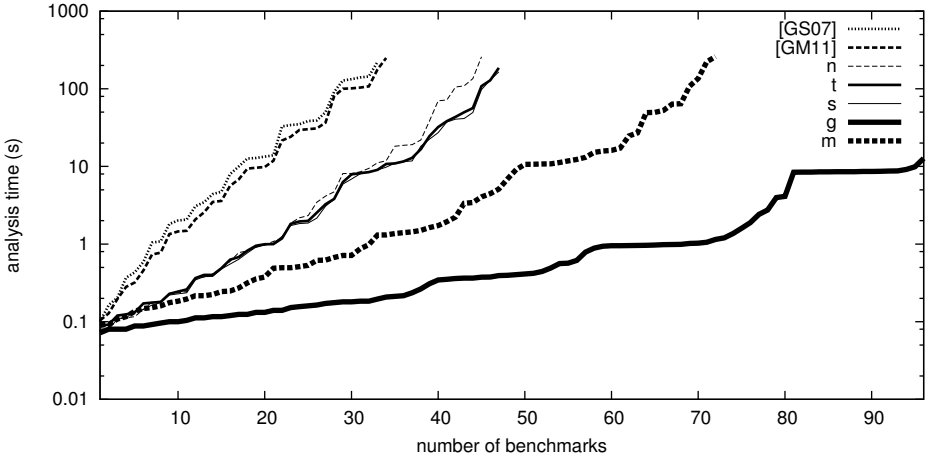

Fig. 1. Comparison of various variants of the max-strategy improvement algorithm. All these algorithms compute the *same* invariant.

(variant (m)) is slower than the variant (g) without the innermost iterations. A possible explanation for this is that as soon as all models have been enumerated, (m) has to confirm unsatisfiability by checking both $F \wedge \neg G$ and $\psi \wedge \neg U$.

However, a broader evaluation is necessary to come to definite conclusions since the structure of the benchmarks is quite simple.

The results of the second comparison (see Table 1) indicate that max-strategy iteration is able to compute better invariants than techniques relying on widening in the *same* $\mathbb{B}^n \to A$ abstract domain. Enhanced widening techniques, such as abstract acceleration, do occasionally improve on precision, but without guarantee to find the best invariant.

We emphasize again that all four max-strategy iteration algorithms in Table 1 compute identical invariants.

An open problem w.r.t. all template-based analysis techniques is however the generation of good templates. For our experiments, we have chosen the weakest of the standard templates (boxes, zones, octagons) that can express the required invariant. Strategy iteration is in general the more expensive technique, but due to our improvements the performance is pushing forward into a reasonable range. These results also show that variant (g) – although a bit slower than (s) in many cases – seems to scale best.

5 Related Work

It has long been recognized that it is a good idea to distinguish states according to Boolean variables or arbitrary predicates (as in *predicate abstraction*). Yet, taking all Boolean variables into account tends to be unbearably expensive: checking the reachability of a state for a purely Boolean program is PSPACE-complete, in practice often solved by constructing a BDD describing reachable

		size						previous algorithms				this paper				std.		abstr.	
		vars				CFG		GS07 [6]		GM11 [15]		g		s		analysis		accel.	
	dom	b	n	bi	ni	lc	ed	time	p	time	p	time	p	time	p	time	p	time	p
Traffic 1	B	6	6	0	0	18	61	2.16	✓	2.10	✓	2.33	✓	2.16	✓	1.22	✓	**0.43**	✓
Traffic 2	Z	6	8	0	0	18	151	122	✓	114	✓	108	✓	**97.0**	✓	3.49		2.86*	
Traffic 3	Z	8	8	1	0	50	619	674	✓	640	✓	357	✓	**329**	✓	22.1		19.2*	
Thermostat 1	B	4	3	0	2	6	15	0.36	✓	0.32	✓	0.28	✓	**0.26**	✓	0.82		0.85	
Thermostat 2	B	6	5	0	4	18	145	16.8	✓	15.1	✓	3.44	✓	**3.23**	✓	26.6		30.4	
Thermostat 3	B	8	7	0	6	66	1357	720	✓	715	✓	66.5	✓	**61.9**	✓	674		908	
Window 1	O	9	5	5	0	21	120	109	✓	102	✓	**70.7**	✓	73.4	✓	4.57		4.70	
Window 2	O	11	5	6	0	45	452	394	✓	372	✓	**189**	✓	286	✓	18.57		23.5	
Window 3	O	13	5	7	0	81	1388	1412	✓	1220	✓	**242**	✓	697	✓	70.2		93.5	
DrugPump 1	B	4	10	4	1	6	231	92.6	✓	90.3	✓	6.05	✓	**4.55**	✓	210		120	
DrugPump 2	B	7	12	8	1	34	11201	timeout > 1800				149	✓	**95.5**	✓	timeout > 1800			
DrugPump 3	B	10	14	8	1	146	112561	timeout > 1800				**1019**	✓	1396	✓	timeout > 1800			

Table 1. Comparison of max-strategy iteration with standard analysis approaches (dom: domain used (boxes (B), zones (Z), octagons (O)); number of variables: Boolean (b), numerical (n), Boolean and numerical inputs (bi, ni); number of locations (lc) and edges (ed) of the control flow graph (CFG); analysis time in seconds; property proved (p); fastest in bold). (* computed with octagons, because zones are not available)

states as the result of a least fixed point computation, which may have exponential size.

While it is possible to encode finite-precision arithmetic into a Boolean program, the large number of Boolean variables and the complicated transition structure generally result in poor performance, thus the incentive to separate arithmetic from "true" Booleans and other small enumerated types. Note that this may not be so obvious for languages such as C or intermediate representations such as LLVM bitcode, where such types may be encoded as integers; a pre-analysis may be necessary [22].

Even if the number n of Boolean variables has been suitably reduced, distinguishing all combinations may be too costly. Various heuristics have therefore been proposed so as to partition $\mathbb{B}^n$ into a reasonably small number of subsets [23]. Relations between the Boolean and numerical states are only kept w.r.t. these equivalence classes [5]. Combining the latter technique with the method presented in this paper to limit partitioning would certainly improve efficiency, however, to the detriment of precision of the obtained invariant which strongly depends on the choice of a clever partitioning heuristics.

Early work in compilation and verification of reactive systems [7] advocated quotienting the Boolean state space according to some form of *concrete* bisimulation. In contrast, we compute coarser equivalences according to per-constraint *abstract* semantics. In the industrial-strength analyzer Astrée, static heuristics determine reasonably small packs of "related" Booleans and numerical variables, such that the values of the numerical variables are analyzed separately for each Boolean valuation [24, §6.2.4]. In contrast, our equivalence classes are computed dynamically and per-constraint.

Disjunctive invariants are related to the partitioning approach; in both cases the invariant is a disjunction $C_1 \vee \cdots \vee C_d$ where C_i are simpler invariants (typically, conjunctions of certain types of literals), but in the disjunctive invariant

approach the C_i may overlap (that is, not have pairwise empty intersection). In such a system, union (as at control merge points) may be implemented by simple concatenation of the disjunctions, but this quickly leads to a blowup; instead a criterion could be used to merge those C_i, *e.g.*, of which the abstract union is actually exact; a similar problem occurs with widening operators [25]. An alternative, which bears some limited resemblance to our strategy-based approach, is to build a map σ meaning that disjunct C_i flows through path π into disjunct $C_{\sigma(i,\pi)}$ [3].

The strategy iteration we have applied proceeds "upward", by successive under-approximations of the least inductive invariant inside the domain converging to it in a finite number of iterations; strategies correspond to paths inside the program, which map to "max" operators in a high-level vision of the problem. There also exists "downward" strategy iteration, where strategies correspond to "min" operators (tests inside the programs and internal reductions of the abstract domain): iterations produce successive *over*-approximations of the least inductive invariant [26, 27], to which convergence is ensured in some cases. A bonus of such an approach is that each iteration produces an over-approximation of the least inductive invariant inside the domain, which may be used to prove safety properties without having to wait for convergence. Sadly, it does not seem to be easily adapted to approaches based on SMT solving, since the SMT formulas would contain universal quantifiers, which greatly complicates their solving.

Recently, a tool for *optimization modulo theory* was presented [14]. We plan to test the variant of our algorithm maximizing Δ (see §3.1) with the help of this tool.

6 Conclusion

We propose a method for computing strongest invariants in linear template domains when the control states are partitioned according to n Booleans or arbitrary predicates, thereby producing a combination of predicate abstraction and template polyhedral abstraction. In accordance with preceding works [1, 3, 4], it traverses loop-free parts of the control graph without need for intermediate abstraction, thus improving the precision. Our method performs strategy iteration, and dynamically partitions the states according to an equivalence relation depending on the current abstraction at each step. The final result is optimal in the sense that it is the strongest invariant in the abstract domain, which a naive algorithm would obtain in at least exponential time and space. While an upper bound on the number of equivalence classes in our algorithm is also exponential n, it can be limited arbitrarily, with some loss of precision. The upper bound on the number of iterations is doubly exponential in n. Our experimental results demonstrate the significant performance impact of various optimizations and the ability to compute more precise invariants in comparison to widening-based techniques.

References

1. Gawlitza, T.M., Monniaux, D.: Invariant generation through strategy iteration in succinctly represented control flow graphs. Logical Methods in Computer Science (2012) Journal version of an article in ESOP 2011.
2. Halbwachs, N., Proy, Y.E., Roumanoff, P.: Verification of real-time systems using linear relation analysis. Formal Methods in System Design **11** (1997) 157–185
3. Henry, J., Monniaux, D., Moy, M.: Succinct representations for abstract interpretation. In: Static analysis (SAS). Volume 7460 of LNCS. (2012) 283–299
4. Monniaux, D., Gonnord, L.: Using bounded model checking to focus fixpoint iterations. In: Static analysis (SAS). Volume 6887 of LNCS. (2011) 369–385
5. Schrammel, P., Subotic, P.: Logico-numerical max-strategy-iteration. In: VMCAI. Volume 7737 of LNCS. (2013) 414–433
6. Gawlitza, T.M., Seidl, H.: Precise relational invariants through strategy iteration. In: Computer Science Logic. Volume 4646 of LNCS. (2007) 23–40
7. Bouajjani, A., Fernandez, J.C., Halbwachs, N., Raymond, P.: Minimal state graph generation. Sci. Comput. Program. **18** (1992) 247–269
8. Miné, A.: Weakly relational numerical abstract domains. PhD thesis, École Polytechnique, Palaiseau, France (2004)
9. Monniaux, D.: Automatic modular abstractions for template numerical constraints. Logical Methods in Computer Science (2010)
10. Gawlitza, T.M., Seidl, H.: Solving systems of rational equations through strategy iteration. ACM Trans. Program. Lang. Syst. **33** (2011) 11:1–11:48
11. Bryant, R.E.: Graph-based algorithms for boolean function manipulation. IEEE Trans. Computers **35** (1986) 677–691
12. Nieuwenhuis, R., Oliveras, A.: On SAT modulo theories and optimization problems. In Biere, A., Gomes, C.P., eds.: SAT. Volume 4121 of LNCS. (2006) 156–169
13. Sebastiani, R., Tomasi, S.: Optimization in SMT with $\mathcal{LA}(\mathbb{Q})$ cost functions. In: IJCAR. Volume 7364 of LNCS. (2012) 484–498
14. Li, Y., Albarghouthi, A., Kincaid, Z., Gurfinkel, A., Chechik, M.: Symbolic optimization with SMT solvers. In: POPL, ACM (2014) 607–618
15. Gawlitza, T.., Monniaux, D.: Improving strategies via SMT solving. In: ESOP. Volume 6602 of LNCS. (2011) 236–255
16. Schrammel, P.: Logico-Numerical Verification Methods for Discrete and Hybrid Systems. PhD thesis, Université de Grenoble (2012)
17. Gonnord, L., Halbwachs, N.: Combining widening and acceleration in linear relation analysis. In: Static analysis (SAS). Volume 4134 of LNCS. (2006) 144–160
18. Schrammel, P., Jeannet, B.: Applying abstract acceleration to (co-)reachability analysis of reactive programs. J. of Symb. Comp. **47** (2012) 1512–1532
19. Bonakdarpour, B., Kulkarni, S.S., Arora, A.: Disassembling real-time fault-tolerant programs. In: EMSOFT, ACM (2008) 169–178
20. Schrammel, P., Melham, T., Kroening, D.: Chaining test cases for reactive system testing. In: ICTSS. Volume 8254 of LNCS. (2013) 133–148
21. Sankaranarayanan, S., Homaei, H., Lewis, C.: Model-based dependability analysis of programmable drug infusion pumps. In: FORMATS. Volume 6919 of LNCS. (2011) 317–334
22. Jeannet, B., Sotin, P.: Inferring effective types for static analysis of C programs. Elec. Notes in Theoretical Comp. Sci. **288** (2012) 37–47
23. Schrammel, P., Jeannet, B.: Logico-numerical abstract acceleration and application to the verification of data-flow programs. In: Static analysis (SAS). (2011) 233–248

24. Blanchet, B., Cousot, P., Cousot, R., Feret, J., Mauborgne, L., Miné, A., Monniaux, D., Rival, X.: A static analyzer for large safety-critical software. In: PLDI, ACM (2003) 196–207
25. Bagnara, R., Hill, P.M., Zaffanella, E.: Widening operators for powerset domains. Int. J. on Software Tools for Technology Transfer **8** (2006) 449–466
26. Gaubert, S., Goubault, E., Taly, A., Zennou, S.: Static analysis by policy iteration on relational domains. In Nicola, R.D., ed.: ESOP. Volume 4421 of LNCS. (2007) 237–252
27. Sotin, P., Jeannet, B., Védrine, F., Goubault, E.: Policy iteration within logico-numerical abstract domains. In: ATVA. Volume 6996 of LNCS. (2011) 290–305